

CLAIMS

- 1 1. A system adapted to distribute redundant information across disks of an array, the
2 system comprising:
 - 3 a storage operating system configured to invoke storage operations executed by a
4 storage system, the storage operating system further configured to manage storage of in-
5 formation, including the redundant information and data, on blocks of the disks in re-
6 sponse to disk access operations, the storage operating system including a storage module
7 adapted to compute the redundant information in response to placement of the data in
8 stripes across the disks, the storage operating system maintaining at least one unallocated
9 block per stripe for use by the storage module to store the computed redundant informa-
10 tion, wherein the at least one unallocated block used to store the redundant information is
11 located in any disk and wherein the location of the at least one unallocated block use to
12 store the redundant information can change over time.
- 1 2. The system of Claim 1 wherein the storage module is a disk array controller config-
2 ured to compute the redundant information and reconstruct blocks lost due to failure of
3 one or more of the disks.
- 1 3. The system of Claim 1 wherein the storage module is a RAID system configured to
2 compute the redundant information and reconstruct blocks lost due to failure of one or
3 more of the disks.
- 1 4. The system of Claim 3 wherein the storage operating system is further configured to
2 implement a high-level module that maintains information about locations of the data on
3 the disks.
- 1 5. The system of Claim 4 wherein the high-level module is a file system or database
2 adapted to control layout of the data on the disks.

- 1 6. The system of Claim 5 wherein the storage operating system integrates the file system
- 2 or database with the RAID system.

- 1 7. The system of Claim 6 wherein the file system or database is configured to determine
- 2 block locations of the data on the disks and the RAID system is configured to determine
- 3 the block locations of the redundant information on the disks.

- 1 8. The system of Claim 6 wherein the file system or database is configured to determine
- 2 block locations of the data and the redundant information on the disks.

- 1 9. The system of Claim 8 wherein the file system or database renders balancing decisions
- 2 to determine the block locations of the data and the redundant information on the disks.

- 1 10. The system of Claim 9 wherein the balancing decisions comprises one of different
- 2 sizes of disks, different speeds of disks, and whether a disk is more heavily utilized than
- 3 other disks.

- 1 11. The system of Claim 8 further comprising block allocation map structures used by
- 2 the file system to determine the block locations of the data and the redundant information
- 3 on the disks.

- 1 12. The system of Claim 11 wherein the redundant information is parity.

- 1 13. The system of Claim 1 wherein the storage operating system comprises computer-
- 2 executable code operable to perform a storage function in the storage system.

- 1 14. The system of Claim 1 wherein the storage module selects the at least one unallocated block to store the redundant information and wherein the storage module computes the redundant information using a redundant storage algorithm.
- 1 15. The system of Claim 14 wherein the selection of the at least one unallocated block to store redundant information is independent of the redundant storage algorithm.
- 1 16. The system of Claim 15 wherein the redundant storage algorithm is a symmetric algorithm or an asymmetric algorithm.
- 1 17. The system of Claim 16 wherein the redundant information is parity.
- 1 18. The system of Claim 14 wherein the at least one unallocated block used to store the redundant information comprises two or more unallocated blocks used to store the redundant information.
- 1 19. The system of Claim 18 wherein the selection of the unallocated blocks to store redundant information is independent of the redundant storage algorithm used to compute the redundant information.
- 1 20. The system of Claim 19 wherein the redundant storage algorithm depends on positions of the blocks in the array.
- 1 21. The system of Claim 20 wherein the redundant storage algorithm is one of a symmetric and asymmetric algorithm and wherein the redundant information is parity.
- 1 22. A method for distributing redundant information across disks of an array, the method comprising the steps of:

- 3 dividing each disk into blocks, the blocks being organized into stripes such that
4 each stripe contains one block from each disk;
- 5 selecting any blocks in the stripe not used to contain data to contain the redundant
6 information, wherein the block used to contain the redundant information is located in
7 any disk and wherein the location of the block used to contain the redundant information
8 can change over time; and
- 9 computing the redundant information based on contents of all other blocks in the
10 stripe, regardless of whether the blocks contain data.
- 1 23. The method of Claim 22 further comprising the step of determining which block in a
2 stripe contains redundant information each time there is a write request to the stripe.
- 1 24. The method of Claim 23 further comprising the step of assigning a block to contain
2 redundant information when each stripe is written.
- 1 25. The method of Claim 24 wherein the step of determining is performed by a high-
2 level module of a storage system and wherein the steps of computing and assigning are
3 performed by a storage module of the storage system.
- 1 26. The method of Claim 25 further comprising the steps of:
2 maintaining, by the high-level module, at least one unallocated block per stripe
3 for use by the storage module; and
4 providing an indication from the high-level module to the storage module of the
5 unallocated block to contain parity.
- 1 27. The method of Claim 26 further comprising the step of reconstructing, using the
2 storage module, a block that is lost due to failure of a disk.

- 1 28. The method of Claim 25 wherein the high-level module is a file system and wherein
- 2 the storage module is one of an array controller and a RAID system.

- 1 29. The method of Claim 28 wherein the step of computing comprises the step of com-
- 2 puting the redundant information in response to placement of the data in stripes across the
- 3 disks.

- 1 30. The method of Claim 29 wherein the step of computing further comprises the step of
- 2 computing the redundant information using algebraic and algorithmic calculations in re-
- 3 sponse to the placement of the data on the array.

- 1 31. Apparatus for distributing redundant information across disks of an array, the appa-
- 2 ratus comprising:
 - 3 means for dividing each disk into stripes, with each stripe containing one block
 - 4 from each disk;
 - 5 means for selecting any blocks in the stripe not used to contain data to contain re-
 - 6 dundant information, wherein the block used to contain the redundant information is lo-
 - 7 cated in any disk and wherein the location of the block used to contain the redundant in-
 - 8 formation can change over time; and
 - 9 means for computing the redundant information based on contents of all other
 - 10 blocks in the stripe, regardless of whether the blocks contain data.

- 1 32. The apparatus of Claim 31 further comprising means for determining which block or
- 2 blocks in a stripe holds redundant information each time there is a write operation to the
- 3 stripe.

- 1 33. A computer readable medium containing executable program instructions for distrib-
- 2 uting parity across disks of an array, the executable instructions comprising one or more
- 3 program instructions for:

- 4 dividing each disk into stripes, with each stripe containing one block from each
- 5 disk;
- 6 selecting any blocks in the stripe not used to contain data to contain parity,
- 7 wherein the block used to contain the parity is located in any disk and wherein the loca-
- 8 tion of the block used to contain the parity can change over time; and
- 9 computing the parity based on contents of all other blocks in the stripe, regardless
- 10 of whether the blocks contain data.